Project-domain Science Traceability and Alignment Framework (P-STAF): Analysis of a Payload Architecture

Laura Jones-Wilson
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
Laura.L.Jones@jpl.nasa.gov

Sara Susca
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
Sara.Susca@jpl.nasa.gov

Kirk Reinholtz
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
William.K.Reinholtz@jpl.nasa.gov

Abstract—**Large science-focused space missions often have multiple instruments working together to address broad science goals. Systems engineers on these types of projects must work with the project scientists to evaluate trades and make decisions that result in a system that efficiently serves the mission science goals. This collaboration is more effective if the systems engineers understand both the traceability from the L1 customer requirements to the selected instruments and the contributions of each instrument in the context of the whole payload suite. These relationships might be understood implicitly by the science team on a project, but there is value in formally codifying them so this understanding can be accessed and formally analyzed by a broader systems engineering effort. We first described a framework for this communication, called the Project-domain Science Traceability and Alignment Framework (P-STAF), in the IEEE 2017 paper "A Framework for Extending the Science Traceability Matrix: Application to the Planned Europa Mission." This paper shows how that basic framework can be leveraged to not only formally capture these relationships between the instruments and the customer needs, but also how that information can be codified in an analyzable graph that can be queried to provide a better understanding of mission risks and scope. This work was drawn from the application of P-STAF to the Europa Clipper mission, but generic example networks are used to illustrate the power of this technique.**

TABLE OF CONTENTS

1. INTRODUCTION

Systems engineers are often tasked with conducting and resolving complex trade studies, which requires making decisions with limited information. For large missions, this complexity is compounded by the fact that relevant information is not always generated or organized by the same people ultimately responsible for those decisions. The Science Traceability and Alignment Framework (STAF) was created to structure the language between scientists and engineers by defining a formal taxonomy and redefining the types of questions that systems engineers ask in the development of the science requirements flowdown. A common language allows both the systems engineers and the project scientists to work together to make informed critical decisions throughout the lifetime of multi-instrument missions.

The taxonomy of STAF was first proposed in previous work ([1],[2]), but in this paper, we explore the evolution of the concept into a graph-based network of information that can be formally queried to gain insight into the science flowdown it represents. This analysis can address questions about the robustness of the payload architecture to failures and faults, help to identify instruments that represent critical nodes to the achievement L1 requirements, and understand the impact that a given L1 or instrument has on the rest of the system.

STAF Basics

The STAF is a taxonomy for formally organizing the science requirements flowdown, and is shown in Figure 1. The first step in its implementation involves decomposing L1 requirements or mission objectives into science "themes," which represent specific science investigations to be done or hypotheses to be tested. (In previous work [2], these were called campaigns but the terminology was changed to avoid overloading the term.) A theme can be investigated using different approaches. For example, on Europa Clipper, the threshold theme Ocean Properties is focused on proving the existence of a subsurface ocean on Europa. There are a number of ways to test the hypothesis that a subsurface ocean exists, including: induction, radio frequency probing, gravity measurements (k_2), obliquity and shape, tidal amplitude (h_2), and libration. Each of these approaches to the Ocean Properties theme might involve different instruments and observations. This approach layer adds specificity and allows a more precise pinpointing of the ways in which the various instrument contributions roll up to the L1 customer requirements. If an instrument contributes to a science theme (via an approach), then a science dataset can be defined using the theme name and the instrument measurement class, such as the Ocean Properties Radar dataset. A science dataset is composed of data taken from a collection of observation types defined by the instrument-specific technique used to collect it and the specific geometric conditions under which it needs to be taken (see Figure 1). An instrument is a separable system designed primarily to collect science observations.

STAF is divided into the project and measurement domains to address the concerns of different stakeholders – those at the project science level and those at the instrument measurement level. The main pivot point between these domains is the science dataset, which is used both by P-STAF (project-domain STAF) and M-STAF (measurement-domain STAF) but in different ways. In M-STAF, the science dataset is the organizing feature that groups the observations and associated measurement requirements by their science contribution. In P-STAF, on the other hand, the role of each science dataset is placed in the context of a particular L1 requirement via the appropriate science theme and approach.

Our previous work [[1] [2]] discusses the development of the STAF in depth, so that material will not be repeated here. However, it is worth recalling that STAF prescribes two main matrix representations of science traceability: 1) an M-STAF matrix associated with each instrument, and 2) and a single project P-STAF matrix. The rows of a specific M-STAF matrix represent the science datasets that a specific instrument contributes to and the observations types needed for those datasets, while the columns specify the qualities that the observations need to have to make them valid for the respective science contributions. The P-STAF matrix, which details how specific instrument and observation types contribute to the L1 requirements, is expanded upon in this work.

Methods of Architecture Evaluation and Graph Analysis

Traditionally, assessments of science robustness, instrument or observation criticality and influence, and the impacts of potential descopes are understood by the science management team in a qualitative sense. The team must leverage their scientific expertise and familiarity with the payload suite to make decisions based on this information and their best judgement. However, few tools exist to help the team manage this information in a formal sense, or to communicate this information across the project. When no systematic way to codify this information exists, incomplete or inaccurate information may be used in the decision-making process instead. While Science Traceability Matrices (STMs) help capture some aspects of this knowledge, they do not always link contributions directly to L1 requirements and may not codify information about the relative importance of the various contributions. [3] For example, when evaluating the payload, one may need to know information like: Which instrument is providing primary data, and which could provide similar data, but less robustly? Which contribution can meet the L1 requirements, and which are enhancing but cannot directly address the requirement? Which contributions are needed together, and which offer independent methods of achieving the science? Traditional STMs do not typically convey this

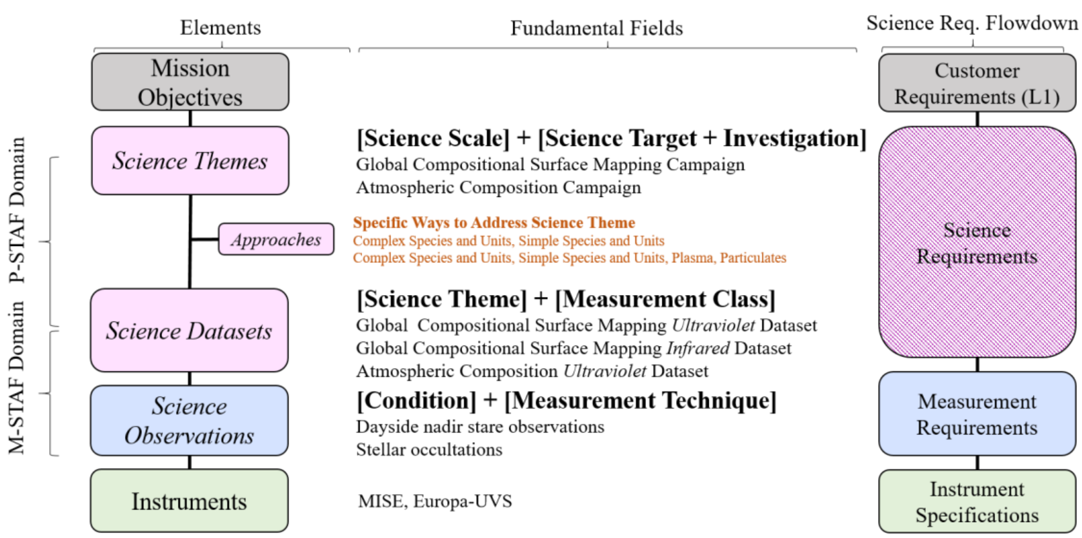

Figure 1 The more mature version of the STAF taxonomy that includes the "Approach" element

kind of information directly. Yet on complicated missions (Europa Clipper is a prime example with almost a dozen investigations), this kind of information needs to be understood and easily communicated across many project stakeholders.

We are proposing an extension of the P-STAF, which organizes the communication between scientists and engineers on the Europa Clipper project, to help the science management team better quantify and track metrics to evaluate a payload architecture, such as robustness in the system and the criticality and influence of specific nodes in the flowdown. To do this, we have developed a modified P-STAF matrix which codifies a network of information similar to a Boolean expression graph (we do not show the low (false) side of the graph). [4] This graph can be analyzed to provide metrics to assess the overall mission architecture. For example, it is possible to identify which combinations of instruments are able to meet a given L1 requirement, to establish which L1 requirements have single point of failures, and to determine how much of the mission is affected by the loss of a particular instrument.

When used together with M-STAF, this P-STAF extension allows for more complex analysis such as determining the amount of resources consumed to meet a project requirement in order to better understand which science objectives require more instruments, observations, data, power, or mass to address. Collectively, this kind of information can quantitatively inform high-level decisions about how to balance risk across the project or how to manage the scope of the mission. Similarly, this information can be used to filter and interpret the results of cross-cutting mission analyses such as those addressing the impact of faults on the likelihood of mission success.

In the sections that follow, we discuss the creation of the P-STAF matrix, how to query the matrix in order to analyze the network of information, and how to interpret those queries into key metrics for each level of the P-STAF hierarchy. We provide examples of both P-STAF networks and analyses based on these networks. The examples are fictitious, but were designed to illustrate key traits and revelations made in the P-STAF network analysis performed for the Europa Clipper mission.

2. DEVELOPING AN ANALYTICAL GRAPH FROM THE P-STAF MATRIX

In the process of developing the P-STAF taxonomy, it became clear that the path between science contributions and the L1 requirements requires a more subtle set of relationships than are often captured in traditional requirement flowdowns. Requirements typically have children that further specify criteria that must be met in order for the parent to be achieved. This relationship implies a fairly rigid "AND" relationship among the children – child requirement A and B and C must be met in order for their parent requirement to be met (to within the margin between the levels). However, in a complex payload where multiple instruments might offer different ways of achieving a given science requirement, there may be multiple paths – with different sets of observations, instruments, and requirements – that could address the science that the project needs. In order to understand the power of this flexibility and the robustness it provides, these "OR" relationships must be identified alongside the traditional "AND" relationships. The P-STAF Matrix [2] can be modified to include this information and thus provide a powerful tool for understanding the payload architecture.

Elements of the P-STAF Matrix

The P-STAF matrix is a tabular way of representing the intersection of the L1 requirements (and associated science themes and approaches) with the contributions of specific observation types and/or instruments. The elements of the P-STAF, i.e., Mission Goal, Mission Objectives (codified in the L1 requirements), Science Themes, and Approaches, are included in columns from left to right. Once each of the rows is defined, the remaining columns are populated with the names of the specific instruments and observation types (the remaining elements of the STAF), as shown in the invented schema in Figure 2.

Relationship Indicators in the P-STAF

Once the matrix has been constructed, it is important to identify the relationships between L1 requirements, the science themes, and the approaches. In this network, we allow for three types of relationships: ENHANCING, AND, and OR. These designations must be assigned and managed by the science leadership team because it requires a scientist's understanding of the interplay between the different elements, and the ability to make a cogent argument for why some approaches or themes are sufficient or not for addressing specific science questions. A small subset of the Europa Clipper mapping is shown in Table 1 as an illustration of how these science relationships may be made.

Table 1 Preliminary L1 requirement example mappings to science themes and approaches for Europa Clipper.

Preliminary Europa Cipper Examples					
Threshold L1 Req.	**Science Theme**	**Theme Definition**	**Can Theme Meet the L1?**	**Approaches**	**Can Approach Meet the Theme?**
Confirm the presence of a subsurface ocean, and constrain whether the ice shell is in a "thin" (several km) or "thick" (10s km) regime.	Ice Shell Properties	Thickness and thermophysical properties of the ice shell.	Yes, together with Ocean Properties	Induction	Yes
				RF Probing	Yes
	Ocean Properties	Existence of the ocean.	Yes, together with Ice Shell Properties	Induction	Yes
				RF Probing	No
				Gravity (K2)	Yes
				Obliquity and Shape	No
				Tidal Amplitude (H2)	Yes
				Libration	No
Search for current activity, notably plumes or thermal anomalies.	Remote Plume Search and Char.	Remote detection of active plumes and their extent above the surface of Europa.	Yes	Volatiles	Yes
				Particulates	Yes
	In-Situ Plume Search and Char.	In-situ detection -of recent or active plumes	Yes	Atmospheric Particulates	Yes
				Atmospheric Volatiles	Yes
				Magnetic Compression	Yes
				Plasma	Yes
	Surface Thermal Anomaly Search	Thermal signatures of current or recent geological activity.	Yes	Thermal Emission	Yes
	Surface Activity Evidence	Surface properties and/or changes indicative of current or recent activity	No	Deposits	Yes
				Surface Changes	Yes

The "enhancing" relationship is used to identify elements which are not required to meet the level above it. For instance, a given L1 requirement might imply many hypotheses to be tested (codified as science themes), but some of those hypotheses may not be critical to meeting the actual text of the L1 requirement. These relevant but not strictly necessary themes are codified as "enhancing." On Europa Clipper, the text of a threshold L1 requirement says to "Search for current activity, notably plumes or thermal anomalies." As shown in Table 1, the "surface activity evidence" theme was deemed not sufficient for meeting the intent of the requirement, which points out thermal anomalies and plumes specifically. However, understanding surface changes can provide great insight into the recent or current activity on the moon, and may inform the thermal anomaly and plume search investigations as well. Thus, to capture the value of performing this science without overstating its ability to address the L1 requirement, we can label this theme as enhancing.

Similarly, there may be approaches that can illuminate aspects of a theme, but cannot directly address the science questions it poses. For example, the Europa Clipper's threshold "Ocean Properties" science theme addresses the existence of the subsurface ocean. Induction is an approach that can fully address these questions – making it a non-enhancing approach. On the other hand, with the current system design, understanding the obliquity and shape of the body can provide context for addressing these questions, but it is not required to have this information to address the goals of the science theme. Thus, "obliquity and shape" is identified an enhancing approach for this mission.

Once the approaches and themes that are enhancing have been identified, the next step is to evaluate the remaining themes and approaches with a one-to-many

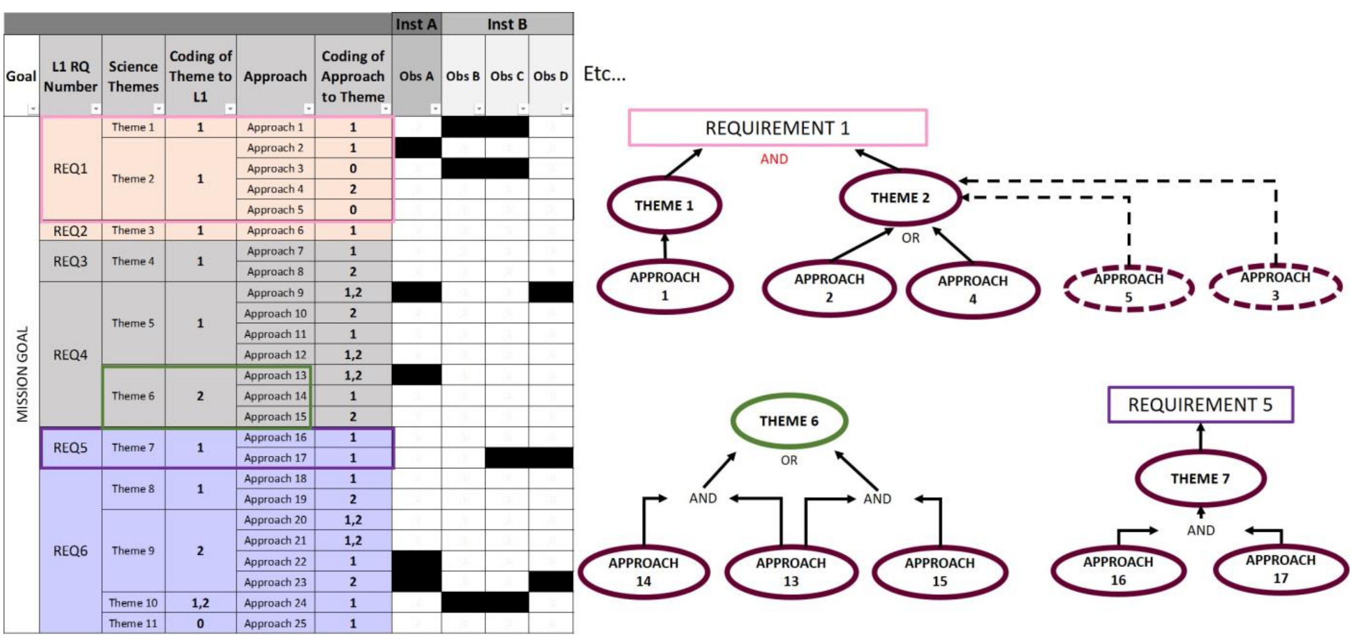

Figure 2 A modification to the P-STAF that enables a graphical understanding of the contributions of each instrument and observation, and provides the framework for further expansion.

relationship. In these cases, it is necessary to identify the "AND" and "OR" relationships between the elements. For example, an L1 requirement for the Europa Clipper Mission (Table 1) states "*Confirm the presence of a subsurface ocean, and constrain whether the ice shell is in a "thin" or "thick" regime.*" This L1 has children themes: "Ice Shell Properties" and "Ocean Properties". Because the L1 requirement explicitly addresses both the Ice Shell and Ocean, these themes have an "AND" relationship, because both themes are needed together in order to address the level above it. On the other hand, the "Remote Plume Search" science theme (Table 1) can be met by either investigating the volatiles or the particulates. Thus, these approaches would have an OR relationship.

In order to easily codify the designation of the AND/OR/ENHANCING relationships into the P-STAF matrix, it is possible to map these relationships to numerical values. In our method, the ENHANCING relationships are numerically codified by a zero, and all other relationships are codified by the rule: "all of group 1 OR all of group 2 OR all of group 3, etc." A theme or approach may fall into multiple groupings, in which case all relevant groups are listed. The numbers are just unique identifiers – they are not ranks or relative valuations. Referring back to our hypothetical P-STAF matrix in Figure 2, the "Coding of Theme to L1" and "Coding of Approach to Theme" columns show how these groups are identified in a P-STAF, and Figure 2 provides examples of how these numerical codes translate into a networked graph.

Contribution Types

Once the full network of AND/OR/ENHANCING relationships has been identified, it is possible to apply this method to the next level down in the STAF: identifying the contributions that science observations make to specific approaches. This step essentially adds information to the dark cells shown in Figure 2. We identify contributions at the observation level because it provides information most directly relevant to the measurement requirements, providing the most clarity. If an observation type addresses a theme or L1 requirement, but does not seem to match to a specific approach, there is a missing approach that should be added to the matrix.

As with the themes and approaches, observation types can be grouped to identify AND/OR/ENHANCING contributions. We maintain the same numerical coding scheme – all observations in group 1 OR all observations in group 2, etc. are identified as ways to achieve the given approach. Enhancing contributions are those in which the data can augment the overall science return, but cannot directly meet the approach as pertinent to a science theme and other non-enhancing contributions do not depend on this data.

It is also useful to further differentiate between the types of non-enhancing contributions. Thus, non-enhancing contributions can be identified as one of three types: primary, independent, and supporting.

1. A primary contribution is one that can provide, most robustly and with greatest probability, the science data necessary to fully achieve a given approach as pertinent to a theme, in the nominal mission plan.

2. An independent contribution is one in which the science data can achieve a given approach as pertinent to a theme, although it may be less robust (e.g., have lower resolution, less coverage, etc.) than primary contributions. Changes to the nominal mission plan (i.e., providing more flyby coverage at lower altitudes) may be required to achieve the approach with this contribution.

3. Supportive contributions identify situations in which science data from one instrument is required to enable contribution of another instrument to fully achieve a given approach as pertinent to a science theme. These support contributions must be in the same coding group as another primary or independent contribution.

Because the contributions may be in multiple coding groups, it is possible but rare that the same contribution is primary with one group and supporting with another; this relationship set is still valid and can still be analyzed. When combined with the "enhancing" contributions, these attributes are referred to as the P/I/E/S designators. Once these designations are assigned to each contribution, and the group numerical code is added, Figure 2 can be transformed to the full P-STAF matrix shown in Figure 3.

Given the STAF taxonomy, the existence of a science dataset (where one instrument contributes to a specific science) is identified by marking the cell belonging to the appropriate row (the approach) and the appropriate column (the observation type) as P, or I, or S, or E. Thus, each marked cell in the P-STAF matrix implies the existence of a row in the M-STAF matrix.

						Inst A	Inst B			Inst C				Inst D	Inst E			Inst F	Inst G	Inst H		Inst I
Goal	L1 RQ Number	Science Themes	Coding of Theme to L1	Approach	Coding of Approach to Theme	Obs A	Obs B	Obs C	Obs D	Obs E	Obs F	Obs G	Obs H	Obs I	Obs J	Obs K	Obs L	Obs M	Obs N	Obs O	Obs P	Obs R
MISSION GOAL	REQ1	Theme 1	1	Approach 1	1		P1	P1					S1									
		Theme 2	1	Approach 2	1	E												P1				
				Approach 3	0		P1	P2					E						E			
				Approach 4	2																	P1
				Approach 5	0					I2												E
	REQ2	Theme 3	1	Approach 6	1							E		P1	E	E				E		
	REQ3	Theme 4	1	Approach 7	1							E		P1						E		
				Approach 8	2							E		P1	I2	I2				E		
	REQ4	Theme 5	1	Approach 9	1,2	P1			E									I2				
				Approach 10	2									P1								
				Approach 11	1									E	I2	I3	I4	P1				
				Approach 12	1,2									E						P1	E	
		Theme 6	2	Approach 13	1,2	P1												I2				
				Approach 14	1									E		P1						
				Approach 15	2									E							P1	
	REQ5	Theme 7	1	Approach 16	1						P1								E			
				Approach 17	1			I2											P1	P1		
	REQ6	Theme 8	1	Approach 18	1										P2	P1	I3					
				Approach 19	2						P1											
		Theme 9	2	Approach 20	1,2															P1	E	
				Approach 21	1,2									I1								
				Approach 22	1	S1												I1				
				Approach 23	2	P1			E									I2				
		Theme 10	1,2	Approach 24	1		E	E			E			I2					P1			
		Theme 11	0	Approach 25	1						P2	P1		I3	I4	I5			I6	E		

Figure 3 A fully coded P-STAF matrix where the P/I/E/S and grouping designators are identified for each observation type as it applies to a specific approach. P = primary, I = independent, S = supportive, E = enhancing

While building the P-STAF contributions mapping, there are a number of points to keep in mind. First, it is important to ensure that the P-STAF only captures contributions that are strong enough to codify as requirements. We emphasize this restriction because the P-STAF analysis will be used to assess robustness with respect to requirements, and if connections that are not captured in requirements are in the P-STAF, the graph will not be consistent with the mission design, invalidating the results. Secondly, it is still valid to capture enhancing contributions to approaches or themes that are not enhancing. For instance, a certain observation type may have a resolution that does not meet the L1 requirement but can still be useful context for the interpretation of the relevant approach/ science theme, even if that science is critical to meeting the L1 requirement. Similarly, it is possible to have a primary or independent contribution to an enhancing theme or approach. It is useful to think about which observation types would make sufficient contributions to the science independent of whether the science is necessary to meet a given L1 requirement. Finally, it is important to note that we do not require a primary contribution for every approach. If an approach has at most an independent contribution identified, the approach itself (given the suite of observation types available) is a less robust way of addressing the theme than approaches with primary contributions. Similarly, if an approach only has at most an enhancing contribution, then the available observation types are not well-suited to addressing that particular approach. Whether this poses a problem is dependent on how the AND/OR/ENHANCING designations in the approach and theme show that approach influencing the L1 requirements. Empty rows or columns simply reveal a mismatch between the science on the mission and the observation types available on the payload. These rows or columns may be useful placeholders, but do not play a role in the analysis discussed later in this paper.

Verifying and Maintaining the P-STAF Matrix

The P-STAF matrix inputs, as with inputs to any analysis, should be validated before being used to inform any decision making process. At least two steps should be performed to ensure that the data is valid:

1) The instrument principal investigators and science teams should be given the opportunity to weigh in on the designations for their observation types and the contributions they make to specific approaches. Where possible, rationales for reconciled contribution designations should be documented and scrutinized by the science team as a whole to ensure data integrity.

2) The P-STAF matrix should be checked against the existing M-STAF matrices, if they exist. If

a contribution in the P-STAF matrix cannot be mapped to a row or set of rows in the M-STAF matrix of the corresponding instrument, then a reconciliation process needs to either add the appropriate measurement requirements to enable that mapping, or removing the contribution from the P-STAF matrix.

If appropriate, an external science peer review of the data may also add a layer of credibility to the data by offering an independent science review.

P-STAF is naturally expressed in tabular format and, the relationships indicators (AND, OR, ENHANCING) as well as the contribution types (Primary, Independent, Supporting, and Enhancing, or P/I/E/S) can be easily expressed using numbers as showed in Figure 3. Hence, a simple way to generate and maintain P-STAF is by using a readily available software such as Microsoft Excel. For the Europa Clipper mission, all the analyses presented in this paper were performed by writing graph manipulation code in MATHEMATICA with the P-STAF Excel file as an input.

Links to the Measurement-Domain STAF

Every identified contribution between an observation type and a science approach implies the existence of (at a minimum) a row in the M-STAF matrix for that instrument. The M-STAF matrix is described in more detail in [1], but essentially each row in M-STAF codifies the set of requirements necessary for a given contribution to be successful.

It is important to note that even enhancing contributions should have corresponding requirements in an M-STAF. Going back to the original mandate for the P-STAF linkages, a contribution should only be captured in this tool if it is strong enough to capture in requirements space. It is true that enhancing contributions do not directly meet an L1 requirement, putting them at risk for descoping, but even so, they are generally a part of the planned mission and do have valuable contributions to make. It is still important to capture that traceability even if these contributions exist outside of the AND/OR relationships that make up the queryable network described in the next section.

3. ANALYSIS FUNDAMENTALS

The full P-STAF matrix as shown in Figure 3 represents a network of information which consists of nodes (L1 requirements, science themes, approaches, science observations, instruments) and directed edges (identified links between the nodes such as an AND relationship). The nodes are simply elements of the STAF and the directed edges are the relationship information captured in the process of building a P-STAF matrix. This network, or graph, and its subsets can be queried to generate a variety of raw analytical products that can be processed into meaningful insights into the graph structure.

Subgraphs can be made from any level in the hierarchy to any other level. These subgraphs can be identified by the two ending levels, or "roots," that they map between. For example, it is useful to parse the data to reveal the relationship between the L1 requirements and the instruments, and so the relevant subgraph will be the one in which the L1 requirement nodes and the instrument nodes are the roots. Because observation types are allowed to contribute to many science datasets, and science datasets are allowed to have many observation types, as described in past work [1] [2], the most insight into the payload architecture comes from understanding relationships between the upper, P-STAF level of the hierarchy (L1 requirements, science themes, or approaches) and the lower, M-STAF level of the hierarchy (science observation types or instruments). Thus, subgraphs with an "upper" and "lower" root are the most powerful tools in this analysis, and will be the focus of the subsequent work.

It is also possible to make subgraphs by removing some of the relationship types. The P-STAF graph codifies all of the P/I/E/S relationships, but it is useful to generate a subgraph (with only primary, independent, and support, or P/I/S, relationships) that removes the enhancing contributions for some assessments because they do not directly impact the system robustness or criticality. In any P/I/S subgraph, the full list of upper-level root nodes are always included in the analysis, even if the node is identified as enhancing when being rolled up to the next level. This approach ensures that when looking at a subgraph of science themes relative to instruments, for example, that all of the science themes appear, even if they are enhancing when they are rolled up to the L1 requirements level.

Using this approach to the P-STAF graph, it is possible to perform a series of queries that, when taken together, can be used to evaluate the strength of the payload architecture as it contributes to the overall mission L1 requirements. There are four basic analyses that can be generated on this graph: the roll-up, the robustness metric, the criticality metric, and the reachability metric.

Roll-Up Analyses

Goal	L1 RQ Number	Science Themes	Coding of Theme to L1	Inst A	Inst B	Inst C	Inst D	Inst E	Inst F	Inst G	Inst H	Inst I
MISSION GOAL	REQ1	Theme 1	1		P	S						
		Theme 2	1	E	E	E			P	E		P
	REQ2	Theme 3	1			E	P	E			E	
	REQ3	Theme 4	1			E	P	I			E	
	REQ4	Theme 5	1	P	E		P	I	P		P	
		Theme 6	2	P			E	P	I		P	
	REQ5	Theme 7	1		I	P				P	P	
	REQ6	Theme 8	1			P		P				
		Theme 9	2	P	E		I		I		P	
		Theme 10	1,2		E	E	I			P		
		Theme 11	0			P	I	I		I	E	

Figure 4 Examples of a roll-up analysis with instrument and science theme at the roots.

The "roll up" analysis is the simplest analysis that can be performed on this graph. In this analysis, each of the contributions is consolidated to show the most "dominant" contribution at a given set of roots. So, for example, for the full P-STAF matrix in Figure 3, the "roll up" analysis with roots of science themes and the instruments would look like the one in Figure 4. In this roll-up, a primary contribution is more dominant than an independent, which is more dominant than a supporting contribution, all of which are more dominant than an enhancing contribution. As can be seen in the figure, a roll-up analysis simply shows the most dominant contribution for each intersection of the roots, and ignores the AND/OR relationships. Enhancing approaches and themes are rolled to the next level by setting all of the contributions in the row to "enhancing." So in the fictitious example in the Figure 3, Instrument C's observation is an independent contribution to approach 5. However, when the roll-up is generated, instrument C has an enhancing contribution to Theme 2 because Approach 5 was enhancing when rolled up to the theme level. If Approach 5 were not enhancing, then the roll up would show Instrument C as an independent instrument for Theme 2. Also note that Theme 11 is enhancing, but the contributions are not set to enhancing yet because this graph does not have an L1 requirement root. A roll-up can also be produced at any combination of L1/theme/approach and instrument/observation.

This roll-up analysis representation allows the detailed information in the main P-STAF analysis to be summarized into the most dominant contributions across the system. This view is useful for providing context for each of the roots. For example, in Figure 4, we can see that Instrument B is really a main contributor to Theme 1 whereas Instrument I primarily contributes to Theme 2. All themes have primary instruments and all instruments make a primary contribution to at least one science theme.

This kind of assessment allows the identification of instruments or observations that are solely enhancing or at most independent (never primary or supporting), which may change the project risk posture towards those elements. For example, in a roll-up with observation type and science themes as the roots, it is possible to see that Observation D is at most enhancing and that Observation E is at most an independent. Both of these observations contribute to the science, but serve different functions than the primary observations, and so systems engineers may choose to avoid letting these kinds of observations drive the system design. If an observation is hard to schedule, then its impact to the P-STAF graph can be queried to understand the effect of that observation on the overall network before expending significant resources to meet it.

However, because the roll-up ignores the AND/OR relationships, it is not possible to tell if Theme 5 is very robust, i.e., can be achieved by either Instrument A, OR D, OR F, OR H or if all of these instruments ANDed

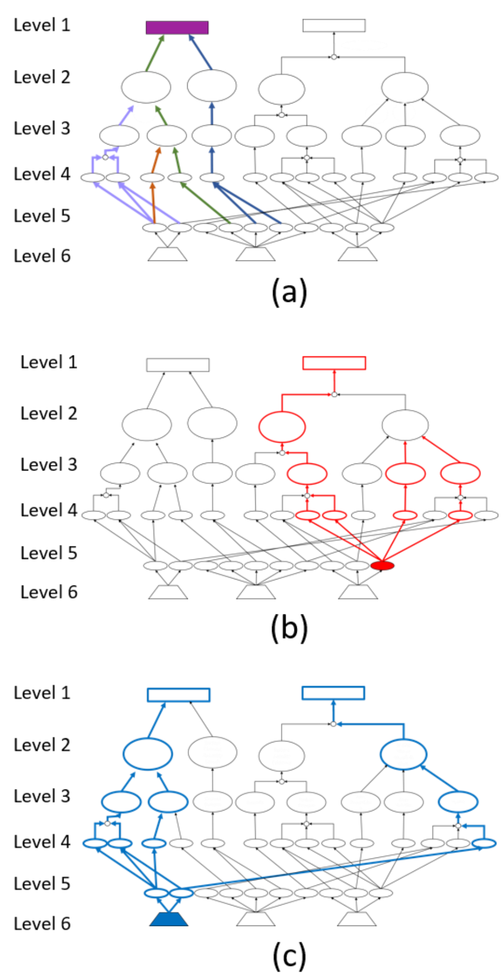

Figure 5 Basic queries of the graph for (a) robustness (b) criticality, and (c) reachability

together are necessary to achieve it. Subsequent graph analyses provide this insight. The power of the roll-up analysis is that it can be performed using the functions readily available in Excel and can thus be a way of quickly checking the decisions made in the building of the P-STAF matrix. It also concisely communicates the role of specific instruments or observations, and helps illustrate the synergies between them.

Truth Tables: Robustness, Minimal Sets, and Scope

Another analysis with high utility is a query of the P/I/S graph to identify all of the combinations of the lower-level root nodes that, when set to TRUE, are necessary to cause an "upper-level" root node to also be TRUE, as shown in Figure 5a. This information can be captured in a "truth table" which lists every minimal combination of TRUE lower-level root nodes that leads to a TRUE in the upper-level root nodes.

For example, the P-STAF matrix in Figure 3 can generate a L1 requirement-to-instrument truth table like that in Table 2. Focusing on REQ1 in Figure 3, we can see that the only way to achieve Theme 1 requires Instruments B (primary) and C (supporting). Theme 2, on the other hand, can be met by Instrument F (via approach 2) or by Instrument I (via Approach 4). Since both Theme 1 and Theme 2 are required to meet REQ1, there are two combinations of instruments that will make REQ1 TRUE: Instruments B, C, and F must all be TRUE, or Instruments B, C, and I must all be TRUE. These two combinations show up in Table 2 as separate rows. In the table, a 0 indicates the instrument is not a part of the combination, and a 1 indicates the instrument is a part of the combination.

Table 2 An example of a truth table created for the notional P-STAF matrix in Figure 3.

REQ to Inst Truth Table	# Inst	Inst A	Inst B	Inst C	Inst D	Inst E	Inst F	Inst G	Inst H	Inst I
REQ1	3	0	1	1	0	0	1	0	0	0
REQ1	3	0	1	1	0	0	0	0	0	1
REQ2	1	0	0	0	1	0	0	0	0	0
REQ3	1	0	0	0	1	0	0	0	0	0
REQ3	1	0	0	0	0	1	0	0	0	0
REQ4	2	1	0	0	0	1	0	0	0	0
REQ4	2	1	0	0	0	0	0	0	1	0
REQ4	2	0	0	0	0	1	1	0	0	0
REQ4	2	0	0	0	0	0	1	0	1	0
REQ5	2	0	1	1	0	0	0	0	0	0
REQ5	3	0	0	1	0	0	0	1	1	0
REQ6	3	1	0	0	1	0	0	0	1	0
REQ6	2	0	0	1	1	0	0	0	0	0
REQ6	2	0	0	1	0	0	0	1	0	0
REQ6	2	0	0	0	1	1	0	0	0	0
REQ6	3	0	0	0	1	0	1	0	1	0
REQ6	2	0	0	0	0	1	0	1	0	0

This kind of analysis reveals information about some aspects of robustness in the system, because the larger the numbers of unique combinations of lower nodes that can achieve an upper node, the more robust the upper node is to failures of the lower nodes. Of course, more robustness can also be interpreted as scope beyond what is strictly necessary to meet a given L1 requirement as well, as will be discussed in Section 4.

Each row in the truth table describes what we call a "minimal set". Because of the way it is constructed, a minimal set is a set of lower-level nodes in which every node is necessary for the success of the upper-level node. In our previous example, for the first row in the truth table, if we only had instruments B, C, and F, the failure of any one of them could cause that instrument combination to no longer meet REQ1. (The same is true for the combination of Instrument B, C, and I). Instrument A or G make contributions to this science, but setting those nodes to TRUE is not required to make REQ1 TRUE; thus, they are not in the minimal set.

It is possible to have different minimal sets for the same upper-level node, as shown in Table 2, where each row represents a different minimal set. Not all minimal sets require the same number of lower-level nodes, as illustrated by REQ5. This requirement (which only has one theme) can be met one of two ways: instruments C and B, or instruments C, G and H. Each of these combinations is a minimal set because, within that combination, each instrument is necessary to make the upper-level node TRUE. However, the first minimal set requires two instruments, while the second requires three. The cardinality of the smallest minimal set expresses the minimum "scope" that the upper-level node requires. The concept of a minimal set and scope is important for explaining the results in Section 4.

Cut Set Tables: Criticality

Another way to query the P/I/S graph is to identify the ways in which setting lower-level root nodes to FALSE can generate FALSE upper-level root nodes, as shown in Figure 5b. This kind of query can identify, for example, how many (and which) L1 requirements fail if a given instrument fails, which illuminates the criticality of a given instrument. This output can also be structured into a table that shows the combinations of failures of the lower-level root that can cause the failure of the upper-level root. We call these tables "cut sets." An example of an L1 requirement-to-instrument cut set table is shown in Table 3.

In cases where it is sufficient for one lower-level FALSE node to cause the upper-level node to also be

FALSE, the lower-level node is called a single-point failure. These single-point failures are sensitive areas in the network where additional resources may be necessary to protect the overall mission, as discussed in Section 4. Single-point failures show up in a cut set table as a row with only a single lower-level node flagged. They also show up in a truth table as an instrument that is flagged in every possible minimal set combination. Combinations of failures (beyond single-point) can provide insight into the effect of multiple descopes, or instruments tied together by common hardware that can fail together.

Table 3 An example of a cut set table created for the notional P-STAF matrix in Figure 3.

REQ to Inst Cut Sets Table	# Inst	Inst A	Inst B	Inst C	Inst D	Inst E	Inst F	Inst G	Inst H	Inst I
REQ1	1	0	1	0	0	0	0	0	0	0
REQ1	1	0	0	1	0	0	0	0	0	0
REQ1	2	0	0	0	0	0	1	0	0	1
REQ2	1	0	0	0	1	0	0	0	0	0
REQ3	2	0	0	0	1	1	0	0	0	0
REQ4	2	1	0	0	0	0	1	0	0	0
REQ4	2	0	0	0	0	1	0	0	1	0
REQ5	1	0	0	1	0	0	0	0	0	0
REQ5	2	0	1	0	0	0	0	1	0	0
REQ5	2	0	1	0	0	0	0	0	1	0
REQ6	4	1	0	1	0	1	1	0	0	0
REQ6	3	0	0	1	1	1	0	0	0	0
REQ6	3	0	0	1	0	1	0	0	1	0
REQ6	2	0	0	0	1	0	0	1	0	0

As an example, consider REQ5 in Figure 3. In order to be TRUE, this requirement needs both Approach 16 and Approach 17 to be TRUE. If Instrument C is FALSE, then Theme 16 is false, and the whole requirement fails. This case shows up as the first row for REQ5 in Table 3, where Instrument C is highlighted. Since there is only one instrument in this row, Instrument C is a single-point failure. Another way for REQ5 to fail is for Approach 17 to fail. Since Approach 17 has two possible ways to be met by different instruments (Instruments B or Instruments G and H), in order to cause REQ5 to be FALSE, the failure must occur in a combination of instruments that removes both of the OR branches. This combination shows up as the second and third row in the example cut set table. If Instrument B and G fail together, or instrument B and H fail together, Approach 17 is unachievable. In this example, instrument C is the most critical because it is a single-point instrument failure for this requirement. Instrument B is the next most critical because it appears in two combinations of double-point failures, making it more critical than Instrument G or H for this particular L1 requirement. This insight, which may not be obvious in the full P-STAF matrix alone, can be used to balance risk, resources, and margin across the payload accordingly.

Reach Tables: Reachability

Finally, a query that can be run on the P/I/E/S or P/I/S graph is one that identifies how many upper-level root nodes can be linked from a given lower-level root node, as shown in Figure 5c. This query, in effect, identifies the link density of a given node to the rest of the network. Thus, this query can be used to assess how connected a given instrument is in the overall network of paths leading to the L1 requirements. This assessment can be codified as a "reach table" which identifies the number of branches (a collection of edges connected head-to-tail) that link one root node to the other root node. When assessed on a P/I/S graph, it identifies the number of non-enhancing branches in which a given node participates, and when assessed on the P/I/E/S graph, it shows the full participation of the node in the entire graph.

An example L1 requirement-to-instrument reach table is shown in Table 4. In the example, for REQ2, instrument D participates in one link between instruments and REQ2 in the P/I/S graph, but since other instruments also make enhancing contributions, the breadth of those contributions is clear in the P/I/E/S graph query, where Instruments C, E, and H also participate in links to the REQ2. For REQ5, there are two paths between the instruments and the L1 requirements on the P/I/S graph: Instrument C and B or Instrument C, G, and H. Because Instrument C shows up in both paths, its part of the reach table identifies its

Table 4 An example of a reach table created for the notional P-STAF matrix in Figure 3 with the P/I/S graph query on top and the P/I/E/S graph query on the bottom.

REQ to Inst Reach Table PIS	# Paths P/I/S	Inst A	Inst B	Inst C	Inst D	Inst E	Inst F	Inst G	Inst H	Inst I
REQ1	3	0	1	1	0	0	1	0	0	1
REQ2	1	0	0	0	1	0	0	0	0	0
REQ3	3	0	0	0	2	1	0	0	0	0
REQ4	17	4	0	0	1	4	5	0	3	0
REQ5	3	0	1	1	0	0	0	1	1	0
REQ6	15	2	0	1	4	3	2	2	2	0

REQ to Inst Reach Table PIES	# Paths P/I/E/S	Inst A	Inst B	Inst C	Inst D	Inst E	Inst F	Inst G	Inst H	Inst I
REQ1	8	1	3	3	0	0	1	1	0	2
REQ2	5	0	0	1	1	2	0	0	1	0
REQ3	7	0	0	2	2	1	0	0	2	0
REQ4	26	4	2	0	6	4	5	0	5	0
REQ5	4	0	1	2	0	0	0	2	1	0
REQ6	31	2	5	5	5	5	2	3	5	0

larger participation in the graph for those roots as compared to Instrument B, G, or H. For the P/I/E/S graph, there is one additional branch: Instrument G offers an enhancing contribution to Approach 16, which appears as an added branch in the P/I/E/S reach table for instrument G.

These results can inform the breadth of influence and reach of a given node, which can be combined with other assessments to understand the impact of a given node.

4. EXAMPLE ANALYTICAL RESULTS

Combining truth tables, cut sets tables, and reach tables allows us to effectively use the P-STAF matrix to gain insight into the structural strengths and weaknesses of various aspects of our payload architecture. For the Europa Clipper mission, once a version of P-STAF is ready, code written with MATHEMATICA takes in the P-STAF excel file and generates the truth tables, cut sets, and reach tables. After that, the summary analyses tables presented in this section can be autonomously generated.

We distinguish between three points of view that use subgraphs with different upper-level root nodes to understand different aspects of the system architecture: the mission level, the L1 requirement or theme level, and the instrument/observation level.

Mission-Level Evaluations

Analyses that use a set of L1 requirements as a root are helpful in providing an overview of the overall P-STAF graph, assessing the robustness of the collection of L1 requirements as a whole, and assessing which lower-level elements are necessary to meet all of the L1 requirements.

There might be more than one minimal set for a group of L1 requirements, and the number of minimal sets and their contents can shed light on the system architecture. The number of minimal sets (for both baseline and threshold mission if specified), can be a metric of the robustness of the L1 requirements. The more possible combinations of lower nodes (i.e., instruments / observations), the more robust the L1 requirements group is. Similarly, the number and type of instruments in each minimal set, can suggest which combinations of instruments are able to achieve all of the given L1 requirements. The intersection of the minimal sets can also be revealing. If there are instruments that do not appear in **any** minimal set, it suggests that these instruments are never single-point failures, and that there is some redundancy in the payload that prevents those instruments from being single-point failures. On the other hand, instruments that appear in **all** minimal sets are those that represent single-point failures for the given collection of L1 requirements.

In projects with both a defined set of baseline and threshold L1 requirements, it is possible to compare the truth tables of the threshold and the baseline L1 requirements to understand how moving between the two sets of requirements changes the graph. In these cases, one might expect that the threshold mission requires a smaller number of instruments in any of its minimal sets when compared to the baseline. Similarly, one might expect to find fewer single-point failures (for both observation type and instrument lower-level roots) in a well-designed threshold set of L1 requirements than in a baseline L1 requirements set.

For the Europa Clipper, the mission-level evaluation tools (truth tables and cut sets) are autonomously generated by in-house graph analysis code. The truth table to the instrument lower-level root is similar to Figure 6 (which is unrelated to Figure 3 and not meant to represent an analysis of the Europa Clipper mission). This example truth table shows the minimal sets necessary for meeting the baseline L1 requirements set and the threshold L1 requirements set. The baseline mission has 8 single point of failures (instruments which appear in all of the baseline minimal sets), while the threshold has none (no instrument appears in all of the 18 minimal sets). Also, note that the threshold

Mission	N of Combos	N of Inst in Combo	Inst A	Inst B	Inst C	Inst D	Inst E	Inst F	Inst G	Inst H	Inst I	Inst J	Inst K	Inst L
Baseline L1s	2	9	1	1	1	1	1	0	1	1	1	1	0	0
		9	1	1	1	1	1	0	0	1	1	1	1	0
Threshold L1s	18	4	0	1	1	0	0	0	1	0	0	0	0	1
		4	0	1	0	1	0	0	1	0	0	0	0	1
		4	1	0	1	0	0	0	1	1	0	0	0	0
		4	1	0	0	1	0	0	1	1	0	0	0	0
		4	0	1	1	0	0	0	1	1	0	0	0	0
		4	0	1	0	1	0	0	1	1	0	0	0	0
		5	0	1	1	0	1	0	0	0	0	0	1	1
		5	0	1	0	1	1	0	0	0	0	0	1	1
		5	0	1	1	0	1	0	0	0	0	1	0	1
		5	0	1	0	1	1	0	0	0	0	1	0	1
		5	1	0	1	0	1	0	0	1	0	0	1	0
		5	1	0	0	1	1	0	0	1	0	0	1	0
		5	1	0	1	0	1	0	0	1	0	1	0	0
		5	1	0	0	1	1	0	0	1	0	1	0	0
		5	0	1	1	0	1	0	0	1	0	0	1	0
		5	0	1	0	1	1	0	0	1	0	0	1	0
		5	0	1	1	0	1	0	0	1	0	1	0	0
		5	0	1	0	1	1	0	0	1	0	1	0	0

Figure 6 Example of a truth table that maps between instruments and set of L1 requirements: the instrument minimal sets needed to meet baseline and threshold L1 requirement sets.

mission has many more minimal sets than the baseline (18 vs 2). This result shows that there are more ways to achieve the threshold L1 requirements (a lower bar) than the baseline L1 requirements. The minimal sets for the threshold mission also require fewer instruments than the baseline, in keeping with the idea that a threshold mission is generally easier to achieve than a baseline.

This truth table also reveals some interesting characteristics of the payload architecture. Instruments C and D, and Instruments J and K are never needed together to meet the L1 threshold requirements, suggesting that, at this level of analysis, they are interchangeable for the threshold mission. Another interesting feature of this analysis is that, Instrument F is never used in any minimal set for the baseline or threshold L1 requirement sets. This result can appear when Instrument F's main contributions can be covered by another instrument, but Instrument F cannot cover all of the science of that other instrument. It is also possible that Instrument F is only enhancing or only contributes to enhancing themes. Further analysis of the graph is necessary to assess the cause of this result. At the mission level, we can only say that if Instrument F was not a part of the payload or suffered some failure that degraded its contributions to the network, the payload would still have all the necessary elements to achieve the L1 requirements. It is possible that some of the science is degraded because an independent instrument must cover for the loss, and the mission may be less robust as a result, subtleties that are covered by additional metrics not discussed in this paper.

	Robustness	Scope	Criticality
	# of Unique Combos of Instruments to Satisfy	Min # of Instruments to Satisfy	# of Instrument Single-Point Failures
REQ1	1	4	4
REQ2	1	6	6
REQ3	1	1	1
REQ4	2	1	0
REQ5	4	3	1
REQ6	1	2	2
REQ7	2	1	0
REQ8	3	1	0
REQ9	10	2	0

Figure 7 Example of robustness, scope, and criticality for a set of L1 requirements.

Level-1 or Science Theme Evaluations

The mission-level analysis offers a birds-eye view of the P-STAF graph, but to appreciate some of the nuances and to better inform trades, it is necessary to dive deeper in the graph. The next level below the mission is the individual L1 requirements and science themes. Using these nodes as the upper-level roots, it is possible to directly compare the L1 requirements or science themes in three ways:

1) Robustness, the number of unique combination of instruments/observations that can satisfy the requirement/theme,

2) Scope: the minimum number of instruments/observations necessary to satisfy the requirement/theme, and

3) Criticality: the number of instruments/ observations that are single point of failures for a requirement/theme.

We found it most useful to use instrument-focused metrics, but for a single instrument mission, a better root would be observation types instead of instruments.

An example of these metrics for a fictitious set of instruments and L1s (unrelated to Figure 3) is in Figure 7. Figure 7, like Figure 6, is similar to the table autonomously generated by the Europa Clipper graph analysis code for the L1 analysis.

The first thing to note is that the L1 requirements can be divided in two groups: those that have some robustness (i.e., can be achieved by more than one combination of instruments) and those that have no robustness and therefore are riskier to meet (i.e., can be achieved by only one combination of instruments). It may be acceptable to have requirements where there is no robustness at the instrument level, depending on the risk posture of the project. In Figure 7: REQ1, REQ2, REQ3, and REQ6 are not robust while REQ4, REQ5, REQ7, REQ8, and REQ9 have some robustness.

To further compare the requirements that have no robustness we have to move to scope and criticality. Note that, if an L1 can be achieved by only one unique combination of instruments, then all the instruments involved in that L1 requirement (scope) are single-point of failures (criticality) by definition. This information can be useful to compare which L1 requirement is probably *more difficult or complex* to meet because it requires more instruments. For example, in Figure 7, REQ1, REQ2, REQ3, and REQ6 are not robust however, it can be inferred that requirement REQ2 is probably the most complex because it needs 6 instruments, while REQ3 can be met by one instrument.

Requirements that have some robustness (REQ4, REQ5, REQ7, REQ8, and REQ9) can still have single-points of failure if all the combinations have instruments in common. As shown in Figure 7, REQ5 has a single-point of failure even though there are four different instrument combinations that can meet it. REQ4, REQ7, REQ8, and REQ9, on the other hand, have some robustness and no single-points of failure at the instrument level, making them the most robust of the L1 requirements in this set. Their minimum scope is also small, as they need only one or two instruments to be met.

This same set of analyses can be performed for the science themes, or against observation-level robustness, scope, and criticality instead. These results allow a direct comparison of the robustness, scope, and criticality of the upper-level root, which can provide feedback to the project leadership team on different levels:

- Do all the L1 requirements/ science themes have single point of failures? How many single points of failure? Is that in line with the risk posture of the mission?

- Do all the L1 requirements/ science themes have a similar scope? Which one is likely to consume more resources than the others?

If the P-STAF graph is paired with the project M-STAF matrices, it is possible to identify the measurement requirements that are single point of failure for each L1 requirement, allowing a deterministic way of identifying the project *key requirements*. Similarly, these tools can help to provide deterministic impact definitions for the classic 5x5 risk matrix for risks related to the measurement requirements and the payload.

Instrument or Observation Type Evaluations

In this last set of analyses, we address the lower-level nodes in the P-STAF graph: the instruments and observation types. The view of the graph from these nodes allows the scientists and engineers to see which instrument has the most impact on the whole mission, and which observation types should be designed with more margin in the instrument design or mission plan.

These analyses focus on the instrument or observation types, turning around the reach and cut set tables to understand the impact of a given instrument on the approaches, themes, and L1 requirements. An example analysis of this type is shown in Figure 8 (fictitious, unrelated to Figure 3, and similar to the table generated for Europa Clipper for instrument level analysis). The first three columns show the instrument reachability for the P/I/S graph, while the next three columns show the reachability including the enhancing paths: P/I/E/S graph. Using this metric, the highest reachability in the P/I/S graph belongs to Instrument C, D, G, and K. Those happen to be also the instruments with the highest reachability in the P/I/E/S graph. This is not always the case. In fact, the instrument reachability in the P/I/E/S graph can be higher or lower for a given instrument than its reachability in the P/I/S graph. If it

	Breadth (P/I/S)			All (+Enhancing) Breadth			Criticality		
Total Number	148 Branches	152 Branches	124 Branches	202 Branches	215 Branches	218 Branches	65 Nodes	29 Nodes	17 Nodes (9BL)
Instrument	% of the Non-Enhancing Branches to **Approach** Influenced by the Instrument	% of the Non-Enhancing Branches to **Theme** Influenced by the Instrument	% of the Non-Enhancing Branches to **L1s** Influenced by the Instrument	% of All Branches to **Approach** Influenced by the Instrument	% of All Branches to **Theme** Influenced by the Instrument	% of All Branches to **L1s** Influenced by the Instrument	% of **Approaches** that Fail if Instrument Fails Alone	Number of **Themes** that Fail if Instrument Fails Alone	Number of **L1s** that Fail if Instrument Fails Alone
Inst A	7%	6%	6%	11%	11%	11%	5%	10%	12%
Inst B	10%	9%	9%	14%	13%	14%	9%	10%	12%
Inst C	18%	16%	15%	25%	24%	24%	9%	14%	18%
Inst D	17%	16%	16%	26%	25%	25%	11%	14%	18%
Inst E	11%	11%	9%	17%	17%	17%	6%	3%	6%
Inst F	5%	5%	3%	10%	9%	10%	0%	0%	0%
Inst G	20%	19%	17%	21%	20%	19%	3%	0%	0%
Inst H	7%	8%	10%	5%	6%	6%	9%	7%	6%
Inst I	7%	8%	10%	6%	7%	6%	6%	7%	6%
Inst J	5%	7%	8%	9%	10%	10%	9%	10%	6%
Inst K	11%	11%	14%	17%	18%	18%	11%	3%	0%
Inst L	1%	1%	2%	3%	3%	3%	3%	0%	0%

Figure 8 Example of reachability, and criticality for a set of instruments.

is higher, it means that that instrument contributes significantly in an enhancing way to approaches/themes/L1s or to approaches/themes that are enhancing. This cannot be said of Instrument H and I. These instruments do not contribute in an enhancing way so, once the E paths are added to the P/I/S graph, their reachability set is the same and, therefore, their percentage values go down.

Although informative, this metric should not be used alone in making decisions in risk and resource allocation. While it is important to know the reachability of an instrument, the information is more telling if it is paired with the instrument criticality (shown in the last 3 columns in Figure 8). These columns identify how many approaches/themes/L1s fail if that instrument fails. For example Instrument G in Figure 8 shows that its reachability on the P/I/S and P/I/E/S graph is high (in the top 4), however, its criticality is quite low. As it turns out, Instrument G has many redundant observation types that offer independent ways of achieving a given set of science. This instrument design means that Instrument G provides robustness to the L1 requirements but it is not a single-point failure. On the other hand, Instrument A and B do not have the highest reachability (only 6th out of 12 instruments), however their criticality for the L1s is only behind the tied Instruments C and D. Hence, it is probably expected that instruments A, B, C and D will be closely monitored to guarantee the success of the L1 requirements that they are critical to achieving. Understanding if an observation or instrument is low impact can help quantitatively guide decisions and mitigate the effects of architectural decisions on the rest of the mission.

5. Extensions and Conclusions

The Science Traceability and Alignment Framework was developed to facilitate communication between scientists and engineers on large science-driven projects such as the Europa Clipper. In its modified form, the P-STAF matrix continues this function by serving as a method of communicating key elements of the science requirements flowdown to a variety of stakeholders (project scientists and systems engineers as well as program manager or review board members). The network codified in the P-STAF matrix also communicates key relationships and synergies between members of the payload, which help illuminate the sensitivities embedded in the payload architecture. On past projects, these sensitivities were generally understood intuitively by project science leadership, there were few available tools for recording and quantifying this understanding. The P-STAF matrix offers a rigorous way to capture this information as a graph so it can be distributed and analyzed more formally. When analyzed with truth tables, cut set tables, and reach tables, the graph can be queried to provide quantitative results to the science team that can inform the science decision-making process. Graph analysis techniques offer ways of evaluating the robustness, criticality, and influence of elements in the payload, and the impact of architectural changes. Ultimately, this tool helps the science leadership team to effectively work with systems engineers to better manage robustness, risk, and resources across the payload by documenting the interplay between elements in the STAF and how the payload contributes to it, exposing this information to verification, external review, and broader accessibility.

The analyses described in this work, however, are not comprehensive. In particular, not all minimal sets are created equal, and differentiators between minimal sets are not described in detail in this work. It is possible to compare minimal sets by their science degradation (the cost of using an independent instrument instead of a primary) or their resource utilization. The P-STAF analysis can also be extended to include its connections to the M-STAF matrices described in [2]. For example, [5] describes how P-STAF and M-STAF can be used to understand fault protection needs. By injecting faults into a nominal mission plan and re-assessing the success of the L1 requirements as filtered through the P-STAF network, it is possible to perform Monte Carlo analyses to get statistical estimates of L1 requirement success. These results can inform risk balancing decisions across observations in a mission plan, allowing mission designers to plan in more observing time for especially critical observation types, for example. Similarly, a mission timeline and the resource expenditure of each observation can be used to compare the data or energy usage of specific L1 requirements.

The full power of this framework is just being brought to bear as we link the P-STAF graph analysis tools to elements of the project that require a complete system-level understanding of the relationships between the mission science and the payload suite. Applying this new method of analysis to the Europa Clipper and other projects will demonstrate its value as a tool to both understand the complex interplays between science instruments and mission requirements, and clearly communicate those relationships to a broad audience across the projects.

ACKNOWLEDGEMENTS

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

The authors would like to thank the Europa Mission science team, project systems engineering team, and mission design team for helping in the development, refinement, and ultimately improvement of the P-STAF concept. In particular, they would like to thank Dr. Robert Pappalardo, Dr. David Senske and Dr. Haje Korth (Europa Clipper Project Scientist and Deputy Project Scientists) for their patience and insightful discussions. In addition, the authors would like to thank the Europa Mission payload office, including Valerie Thomas, Valerie Duval, and Kari Lewis for their enthusiasm and support of this concept.

REFERENCES

[1] S. Susca and L. Jones-Wilson, "A Framework for Writing Science Measurement Requirements and its Application to the Europa Multipl Flyby Mission," in *IEEE Aerospace Conference*, Big Sky, MT, 2017.

[2] L. L. Jones-Wilson and S. Susca, "A Framework for Extending the Science Traceability Matrix and its Application to the Planned Europa Mission," in *2017 IEEE Aerospace Conference* , Big Sky, MT, March 4, 2017.

[3] J. R. Weiss, W. D. Smythe and W. Lu, "Science Traceability," in *IEEE Aerospace Conference Proceedings*, April 2005.

[4] H. H. Henrik Reif Andersen, "Boolean Expression Diagrams," *Information and Computation*, vol. 179, p. 194–212, 2002.

[5] K. McCoy, B. Nairouz, B. Bradley, L. Jones-Wilson and S. Susca, "Assessing Science Robustness of the Europa Clipper Mission: Science Sensitivity Model," in *IEEE Aerospace Conference*, Big Sky, 2018.

BIOGRAPHY

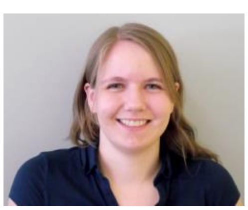

Laura Jones-Wilson *earned her B.S. in Aerospace Engineering at Virginia Tech and went on to study at Cornell University, where she obtained her M.S. and Ph.D. in Aerospace Engineering specializing in Dynamics and Controls. She then joined JPL in 2012 as a guidance and control systems engineer, where she served as the project systems engineer for STABLE. She now serves as the PI a Mars Sample Capture technology development effort, the co-manager of the SmallSat Dynamics Testbed, and the instrument engineer for the radar instrument as a member of the Europa Mission payload team*

Sara Susca *received her PhD in Electrical Engineering from UCSB in 2007. She spent three years at Honeywell Aerospace developing new technology for GPS denied navigation. She has been with JPL since 2011 where she covered various roles including Project Manager for STABLE (a balloon-borne sub-arcsecond pointing demonstration). She is currently member of the Europa Mission Payload team as the instrument engineer for the EIS instrument.*

Kirk Reinholtz *is a Principal Engineer at California Institute of Technology/Jet Propulsion Laboratory, joining in 1989. His software has been to Mars twice, on Mars Observer and Mars Global Surveyor. He joined the NASA ground system program in 2009 to help revitalize the overall software system and architecture. He joined the Europa project in 2014 and has been there ever since. Along the way his research work was awarded the Dukes Choice award in 2004. Prior to joining JPL he worked in the software tech arena, and has an educational software co-publication with Isaac Asimov ("Science Adventure: Discoveries that Changed the World"). He has an MSCS from University of Southern California.*